

What is Claimed:

1. A virtual data center implemented on hardware resources of at least one host computer having at least one host processor and system resources including memory divided into most privileged system memory and less privileged user memory, the virtual data center comprising:

virtualization software loaded on each host computer, said virtualization software operating in said less privileged user memory and dividing the host computer into a plurality of virtual partitions including at least one user guest partition and at least one system partition, said at least one user guest partition providing a virtualization environment for at least one guest operating system, and said at least one system partition maintaining a resource database for use in managing use of said at least one host processor and said system resources;

at least one monitor that operates in said most privileged system memory and maintains guest applications in said at least one guest partition within memory space allocated by said at least one system partition to said at least one guest partition;

a context switch between said at least one monitor and said respective guest and system partitions for controlling multitask processing of software in said partitions on said at least one host processor; and

a software application that owns a configuration policy for said data center and tracks persistence for respective domains to which each partition of said at least one host computer is assigned by said at least one system partition.

2. The virtual data center of claim 1, wherein said at least one system partition comprises an ultravisor partition that includes said resource database and a resource management software application that assigns system resources to respective system and guest partitions and provides an index to the assigned system resource in said resource database, a command partition that owns a resource allocation policy for the host system on which it is loaded and that creates transactions that pass through a command memory channel between said command partition and said ultravisor partition for processing by said resource management software for reallocation of said system resources as specified in said transaction.

3. The virtual data center of claim 2, wherein the software application is implemented in a system partition of one of said at least one host computers and the software application exchanges resource transactions with said command partition.

4. The virtual data center of claim 3, wherein the software application maintains a persistent database of virtual partition definitions for at least one domain of said at least one host computer.
5. The virtual data center of claim 4, wherein the command partition stores a copy of the virtual partition definitions for said at least one domain for bootstrap purposes in the event of a partition failure or a hardware failure of a host computer.
6. The virtual data center of claim 2, wherein upon activation of a partition, the software application selects a host computer of said at least one host computer having required resources for said activated partition, connects to a resource service running in a command partition of said host computer, and provides a definition of the activated partition and a start command to the resource service.
7. The virtual data center of claim 6, wherein said command partition stores a copy of said resource database, uses said copy of said resource database to select appropriate resources for the activated partition, and creates a transaction to update said resource database via said command memory channel.
8. The virtual data center of claim 1, wherein said software application monitors operation of said at least one host computer and, upon detection of host computer failure, chooses a new host computer for virtual partitions assigned to a failed host computer.
9. The virtual data center of claim 1, wherein said software application assigns an interconnected set of system resources of said at least one host computer to a zone and respective partitions are assigned to the zones with the system resources required by the respective partitions, where a zone is unit of resource allocation for system resources of said at least one host computer within a computer network.
10. The virtual data center of claim 9, wherein said software application assigns new partitions to a host computer that does not include said software application by sending, over a network connection, a resource transaction to a command partition of the host computer that does not include said software application.
11. The virtual data center of claim 10, wherein said software application enables migration of an active partition on a first host computer to a second host computer by transferring memory

contents of the active partition from the first host computer to a target partition activated in the second host computer via said network connection.

12. The virtual data center of claim 1, wherein said configuration policy targets allocation of system resources to a zone based on at least one of quality of service guarantees, bandwidth guarantees, and physical location of respective host computers.

13. The virtual data center of claim 1, wherein said configuration policy is changeable by a user to permit changes in configuration of said system resources based on different system resource schedules at different times.

14. The virtual data center of claim 3, wherein a redundant software application is loaded in a system partition in a second host computer different from the host computer hosting said software application.

15. A method of implementing a virtual data center on hardware resources of at least one host computer having at least one host processor and system resources, comprising the steps of:
dividing said at least one host computer into a plurality of virtual partitions including at least one user guest partition and at least one system partition, said at least one user guest partition providing a virtualization environment for at least one guest operating system, and said at least one system partition maintaining a resource database for use in managing use of said at least one host processor and said system resources;

maintaining guest applications in said at least one guest partition within memory space allocated by said at least one system partition to said at least one guest partition;

providing a context switch between said respective guest and system partitions for controlling multitask processing of software in said partitions on said at least one host processor; and

assigning each partition to a domain of said at least one host processor in accordance with a configuration policy for said data center; and

tracking persistence for respective domains to which each partition of said at least one host computer is assigned by said at least one system partition.

16. The method of claim 15, wherein said at least one system partition comprises an ultravisor partition that includes said resource database and said resource management software application and a command partition that owns a resource allocation policy for said host system,

further comprising the steps of creating transactions that pass through a command memory channel between said command partition and said ultravisor partition and said resource management software processing said transaction for reallocation of said system resources as specified in said transaction.

17. The method of claim 16, wherein the steps of assigning system resources comprises the step of assigning each partition of said at least one host computer to a domain based on a configuration policy.

18. The method of claim 17, further comprising the step of maintaining a persistent database of virtual partition definitions for at least one domain of said at least one host computer.

19. The method of claim 18, further comprising the step of storing a copy of virtual partition definitions for said at least one domain for bootstrap purposes for initial startup and in the event of a partition failure or a hardware failure of a host computer.

20. The method of claim 18, wherein upon activation of a partition, performing the steps of selecting a host computer of said at least one host computer having required resources for said activated partition, connecting to a resource service running in a command partition of said host computer, and providing a definition of the activated partition and a start command to the resource service.

21. The method of claim 20, further comprising the step of storing a copy of said resource database in said command partition, using said copy of said resource database to select appropriate resources for the activated partition, and creating a transaction to update said resource database via said command memory channel.

22. The method of claim 16, further comprising the steps of monitoring operation of said at least one host computer and, upon detection of host computer failure, choosing a new host computer for virtual partitions assigned to a failed host computer.

23. The method of claim 18, further comprising the steps of assigning an interconnected set of system resources of said at least one host computer to a zone and assigning respective partitions to the zones with the system resources required by the respective partitions, where a

zone is unit of resource allocation for system resources of said at least one host computer within a computer network.

24. The method of claim 23, wherein said partitions assigning step comprises the steps of assigning new partitions to a host computer by sending, over a network connection, a resource transaction to a command partition of the host computer that is to host the new partition.

25. The method of claim 24, further comprising the step of migrating an active partition on a first host computer to a second host computer by transferring memory contents of the active partition from the first host computer to a target partition activated in the second host computer via said network connection.

26. The method of claim 17, further comprising the step of changing said configuration policy based on different system resource schedules at different times.

27. The method of claim 16, further comprising the step of separating resource management functionality into management by applications in at least three separate partitions, an operations partition that maintains resource allocation policy, a command partition that generates transactions requesting resources in accordance with the resource allocation policy, and a resource management partition that processes the transactions and updates the resource database based on said transaction processing.